SCIENCE INFORMATION SYSTEMS NEWSLETTER

Office of Space Science

Science Information Systems Program

Read the latest on NASA's information systems

www-sisn.jpl.nasa.gov

46

Welcome to Issue 46

1998

Features

The NASA Science Information Systems Newsletter (SISN) is prepared for the Office of Space Science (OSS), Science Information Systems (SIS) Program through an agreement with the Jet Propulsion Laboratory. The newsletter, which has been an ongoing task for over ten years, is a forum for the space science and applications research community to report research and development activities, outreach activities, and technology transference. SISN offers a venue for articles that are not likely to appear elsewhere and provides the opportunity for information exchange within the science community, as well as a platform for accomplishments by that community. Related articles from other programs and agencies are also published.

Questions or comments regarding this newsletter task may be emailed to Sandi Beck at <sandi.beck@jpl.nasa.gov>.

Credits: Design and format by Editor, Sandi Beck. Graphics design, animation, and rollovers by Scott Brenneisen, XTREME GRAFX. Also, thanks over the years to Calvin Yee, formerly of Telos Information Systems, Doug Steinwand, formerly of JPL, and the JPL documentation staff.

Copyright: If use is made of any of this publication's graphics, logos, or icons, please give credit to the NASA Science Information Systems Newsletter.

The Applied Information Systems Research Program (AISRP) maintains an awareness of emerging technologies applicable to space science disciplines, supports applied research in computer and information systems science and technology to enhance NASA Office of Space Science (OSS) programs, stimulates application development where warranted, and provides for the systems analysis and engineering required to transfer new technology into evolving OSS space science programs through NASA Research Announcements.

Improving Image Access Using Progressive Data Compression

Aaron Kiely and Adina Matache, Communications Systems and Research Section, Jet Propulsion Laboratory

With the advent of digital images, scientists and engineers have discovered that not only is a picture worth a thousand words, it also usually requires many thousands of bytes to store or transmit. This presents significant problems in the archiving, management, and dissemination of scientific digital images obtained by NASA space exploration missions. Problems of access, transmission, and retrieval time are becoming increasingly apparent with the popularity of viewing images using the World Wide Web, especially over constrained communications channels.

Image compression techniques provide some relief from the problems associated with huge databases. Popular techniques include the Graphic Interchange Format (GIF) and Joint Photographic Experts Group (JPEG) methods, which are in widespread use on the Web. While these methods can serve as effective tools, both allow significant room for improvement. The GIF engine is a string-matching algorithm that is suitable for some types of computer generated images but generally offers poor performance for digital photographs and scientific imagery. The JPEG technique is better suited to digital photographs, but at high compression ratios it suffers from blocking artifacts endemic to the Discrete Cosine Transform (DCT) on which it is based. Moreover, it is becoming accepted in the data compression community that DCT-based techniques are outperformed for overall image compression efficiency by newer techniques based on wavelet decompositions.

Progressive transmission

While compression reduces the overall time to access and retrieve an image from a remote database, access can be made more efficient by exploiting progressive transmission within the compression algorithm. An image data compression system that uses progressive transmission allows you to reconstruct successively higher fidelity versions of an image as data is received. Segments of compressed data are ordered so that low-and medium-resolution versions of the overall image can be viewed even when only a small fraction of the compressed data has been retrieved. The goal is not just efficient overall compression, but also to provide the highest fidelity image possible at each stage of transmission.

Progressive transmission makes more efficient use of a communications channel, because images become visible more quickly, reducing the apparent retrieval time. Moreover, progressive transmission enables efficient browsing of image databases; based on a quickly obtained preview of an image, you might decide to terminate the transmission and proceed to the next image. Alternatively, you might select a low or medium resolution default for image retrieval. In this way, images of less value are not transmitted at full resolution. Because of the small volume of data required to transmit relatively high quality previews of an image, the savings in access time can be significant when browsing large databases of images.

As part of the Progressive Wavelet Image Compression project, software is being developed to perform progressive image compression and decompression of scientific images based on wavelet decompositions of images. Wavelets have a tremendous advantage for progressive transmission because, in addition to superior overall compression efficiency, wavelet decompositions naturally represent image data in a hierarchical manner. Consequently, wavelet decompositions are well suited to progressive image compression.

Wavelet transforms

Transform-based compression methods such as wavelets decompose the original data before compression. The motivation for such a method is that the data may be more easily compressible after application of a well-selected transform. As a simple example, when pixel intensities in an image are slowly changing, instead of encoding each pixel independently, it may be more efficient to encode the difference between adjacent pixels.

The Discrete Cosine Transform (DCT) is an example of a block-oriented transform technique that is used in the JPEG algorithm. The image is segmented into small blocks, usually 8 x 8 pixels in each block, and each block is transformed before compression. In contrast, wavelet decompositions transform the entire image. An image decomposition using wavelets is illustrated in the following Figure1.

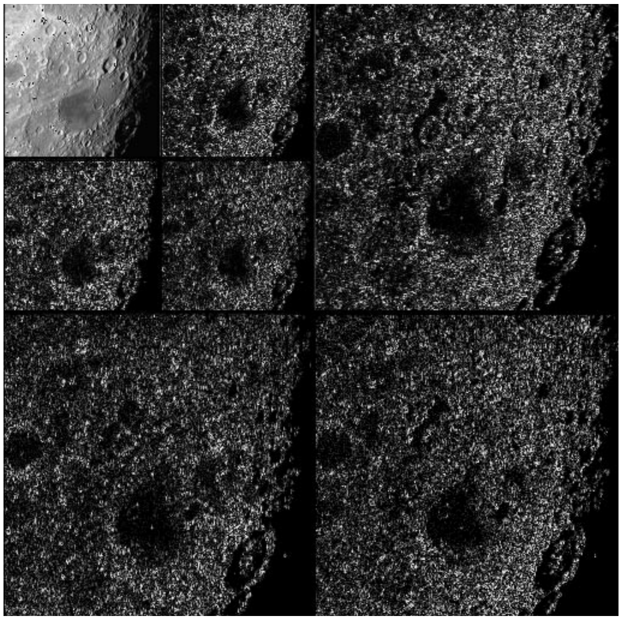

Figure 1. An image decomposed, using wavelets, into "low frequency" (upper left corner) and "high frequency" information.)

A wavelet decomposition of an image is naturally hierarchical. Because energy in images is more concentrated in the low frequencies, by transmitting only the lowest frequency components, we can obtain low resolution previews from a small fraction of the data. For example, in the transformed image Figure 1, note that one of the transformed components (shown in the upper left corner) naturally provides a low-resolution preview of the image.

As other researchers have shown, by carefuly modifying the wavelet decomposition the overall transform is lossless. This means that when all of the compressed data are transmitted, the original image can be exactly reconstructed. The result is that the same algorithm can be used for lossy and lossless compression.

Layering the transformed data

Transmitting the decomposed data in order from lowest to highest frequency is a simple method of incorporating progressive transmission in an image compression scheme. Moderate quality images can be reconstructed even when only the lowest frequency information has been received. In addition to exploiting the hierarchical wavelet decomposition, each subband can be further separated into layers so that the data in an individual subband can be transmitted progressively. This added dimension allows high quality image previews more quickly, while still maintaining high overall compression, because we can switch between the subbands during transmission.

The combination of wavelet coding and layered compression of subband data allows higher quality images to be reconstructed more quickly. At each step in transmission, we would like to transmit the compressed subband layer that gives the largest improvement in image quality per compressed bit. The process of exhaustively determining the optimal ordering of layers is computationally intensive. However, analysis of image subband statistics has yielded a simpler heuristic rule whose performance is close to optimal. This method adapts to the image being compressed, changing the order of transmission depending on the image data encountered, as shown in Figure 2.

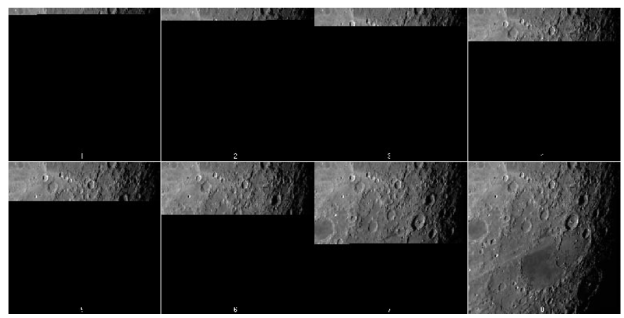

Figure 2. This sequence shows the intermediate images obtained in the retrieval of an image compressed using a raster-scan oriented method

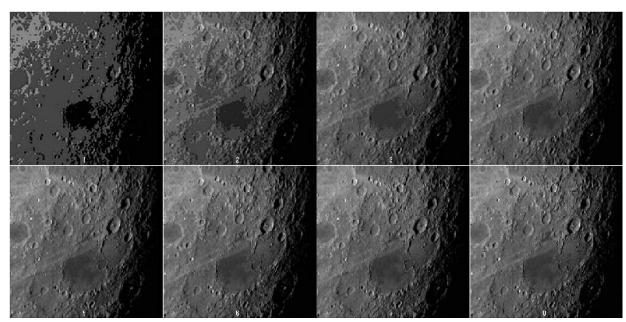

This sequence shows the intermediate images obtained using a progressive wavelet method. For each pair of images with the same number, the same amount of data has been transmitted using each technique.

To learn more about this research, email Aaron Kiely at <aaron@shannon.jpl.nasa.gov>. ■

Focusing radar's lens -

Remote Sensing Tools Reveal Ecological Secrets

Jarrett Cohen, Science Writer, Goddard Space Flight Center

As a young mountain climber, Jeff Dozier witnessed several avalanches safely from above. Their naked power drew him to the science of snow.

"I learned that there are many interesting questions about water and vapor transport in snow, and about snow's electromagnetic 'signature.' I combined my hobby and my job," - Jeff Dozier, Dean of the School of Environmental Science and Management at the University of California, Santa Barbara (UCSB).

Of greatest concern to Dozier and fellow Californians is how much water melted snow will produce in the spring; mountain runoff supplies the bulk of water for agriculture, homes, and work places. The one method to reliably predict this "snow water equivalence" is synthetic aperture radar, or SAR.

Kelly Elder (standing, now at Colorado State University) and the University of California-Santa Barbara's (UCSB) Michael Colee extract and weigh snow cores to predict how much water regional snow will produce when melted.

Synthetic aperture radar "is the only remote sensing technology that sees down into the snow and is the only way to measure the amount of snow" across large areas. - Jeff Dozier, UCSB.

Aircraft, the space shuttle, and satellites have been the mounts for SAR antennae, which bounce radar signals off the Earth and use the return time to gauge the distance. Radar's self-illumination and long wavelengths make SAR an any-time, all-weather instrument. For Dozier, SAR is the only remote sensing technology that sees down into the snow and is the only way to measure the amount of snow" across large areas.

Jeff Dozier and Jiancheng Shi are among the UCSB scientists analyzing world-wide snow scenes captured by NASA's Shuttle Imaging Radar-C (SIR-C).

However singular its traits, SAR is useless without computers. "The signals look like white noise because the radar doesn't have a lens," said Grand Challenge investigator, Dave Curkendall of NASA's Jet Propulsion Laboratory (JPL). "Like the aperture controlling how much light enters a camera lens, SAR gets a wonderful aperture from the motion of the spacecraft. The computer then supplies the lens that focuses the image."

Curkendall estimates that it requires 800 billion floating-point operations (flops) to "focus" one image taken by NASA's SIR-C, for example. Processing regional data like that for snow water equivalence becomes a Grand Challenge problem and thus a task for today's fastest supercomputers. To tame this "computational bear," Curkendall's team is exploiting SAR's unique characteristics for three crucial ecological problems.

Melted snow forecasts

According to Dozier, yearly SAR data processed with the UCSB algorithms could lead to better snow-melt water runoff forecasts. High-performance computing enables widespread measurement of snow water equivalence, the major unsolved problem in snow hydrology. The goal is to process all the snow scenes from SIR-C's 10-day flights of April and October, 1994, months when melting occurs. Validating SIR-C predictions with 1994 runoff data will show what could be done with a permanent-orbiting SAR in space.

Mountainous areas such as the Sierra Nevada, Alps, Himalayas, and Andes are the primary interest for the validation. Because there will be hundreds of scenes involving trillions of bytes, UCSB scientists are adapting their workstation-conceived algorithm for parallel computers.

The snow algorithm transforms raw SIR-C signals into images that are analyzed to determine which areas are snow-covered, the quantity of snow (usually expressed by snow water equivalence) and if snow is wet or dry (wet snow adds to runoff). Combining these factors with weather forecasts pinpoints snow water equivalence down to as accurate as 20 centimeters. Field measurements of the same variable can only cover a few sites in any given time periods. Today, researchers ski across mountains to extract cores and weigh snow by hand and then compare the measurements with river runoff data.

"It gives a pretty good estimate, but like all empirical methods, it occasionally goes wrong, with errors of 30 percent in some years," Dozier said.

“Taking advantage of SAR's ability to look at widely independent regions, the snow algorithm calculates all possibilities”, explained Jiancheng Shi, research assistant at UCSB. “Previously, local conditions limited the measurements. With a US SAR satellite and the snow algorithm's maturation, better runoff forecasts could become reality. In California alone, a one percent improvement would be worth millions of dollars.”

Jiancheng Shi, research assistant at UCSB.

"We are working with the California Department of Water Resources to translate research results into operational practice," Dozier said.

Cracking earthquakes open

As with the water supply's ebbs and flows, a worrisome ecological threat for California is earthquakes. SAR's penetration of Earth's surface and high spatial resolution are proving key to studying tectonic plate movements in the state's southern portion.

Yehuda Bock, Scripps Institution of Oceanography, explains the mechanics of a Global Positioning Satellite (GPS) receiving station, which can monitor tectonic plate movements as small as one millimeter.

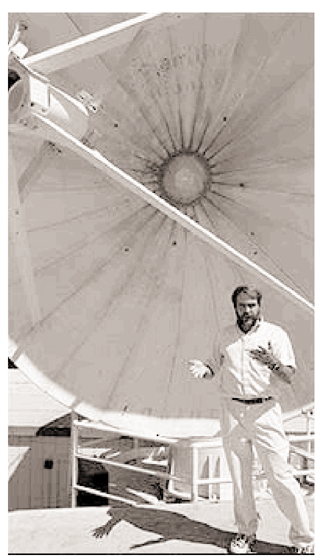

Capturing movement demands repeat-pass SAR interferometry, which correlates data from multiple satellite orbits.

Scripps researchers are focusing on the European Earth Resource Satellites (ERS-1 and ERS-2), for earthquake studies, correlating data from multiple orbits to create interferograms that span tectonic events.

"You take two images that span a tectonic event and fuse them into an interferogram," said David Sandwell, San Diego Institution of Oceanography professor of marine geophysics. "The satellite doesn't always come back to the

same position, adding a parallax effect. What you see is the topography; you have to remove the topograhic signature."

Peeling away that signature hopefully will open up a sweeping view of Southern California deformation every 35 days, Sandwell stated. A complementary network of 50 GPS satellite stations has disclosed total movement along the 400 kilometer plate boundary to be 45 to 50 millimeters per year, about the rate at which your fingernails grow, according to Curkendall.

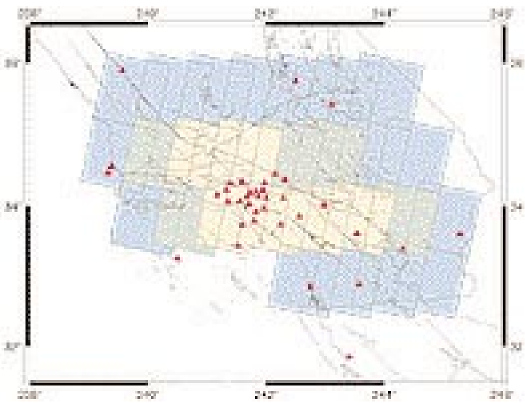

SAR (100-by-100 kilometer squares) and GPS satellite stations (shown by triangles) complement each other in providing a comprehensive view of Southern California plate movement. Lighter shaded areas show increasing overlap density between SAR and 250 proposed GPS stations (now 45).

"SAR gives us the possibility of determining the deformation over the continuum of Southern California, not just a selection of points, but a spatially dense picture," said Yehuda Bock, research geodesist at Scripps. "The reason to combine the two is that these GPS positions are known to a millimeter or so. We use them as calibration points, so the map of deformation is more accurate."

What GPS and earlier approaches cannot do is "get a handle on inhomogeneity in the crustal deformation, including localized effects that may be due to non-tectonic sources, such as water withdrawal," Scripps professor of geophysics,Jean-Bernard Minster, explained. "SAR's regional snapshot can give us a detailed understanding of how the Earth prepares for earthquakes."

Jean-Benard Minster, Scripps Institution of Oceanography, San Diego, CA

With no detection of a particular tremor signal, this knowledge will not likely lead to earthquake prediction. However, "--a fault may be sliding and relieving its stress," Sandwell said, pointing to safe areas for buildings and roads. "Predicting the absence of a big earthquake is almost as important as attempting to predict the earthquake location and time."

Rojoicing at floods

Much more relevant than prediction in the Amazon rainforest are regular and timely surveys showing inundation (or flooding) and deforestation zones. SAR again is uniquely qualified for the duty, as Brazilian farmers practice "slash and burn" agriculture.

"They fell the larger trees and burn them for crop fertilizer," said Tony Freeman, a radar instrument manager at JPL. "The land is completely obscured by smoke, but radar sees right through it."

Tony Freeman, JPL

Japan's Earth Resources Satellite-1 (JERS-1) can cover the United States-sized region in two months, encompassing the flood (April-June) and dry seasons (October-November) that JPL and others are mapping (Freeman and collaborators have completed a JERS-1 map of 1995's dry season). The satellite took 1500 70-square kilometer frames of data. JPL-developed software, called the Digital Light Table, pieced together the frames into a seamless mosaic.

When processed, "there is some overlap between image scenes, much like when you take landscape portraits," Freeman said. "We use correlation matching to determine what that overlap is, then run a large matrix inversion program to co-register images."

The dry season map clearly traces the shape and size of the river course, which had been slightly unknown, explained Freeman. "We can see what areas are inundated during this season across the whole Amazon; this is the primary product. The secondary product is a map of the forest, deforestation, and land cover," he said. Classification studies will calculate these categories every 100 meters.

It takes several years to establish deforestation from optical satellite photos, Freeman reflected. Using SAR, there is the ability to monitor this phenomenon and floods every six months, leading to adjustments in local strategies.

"When the Amazon floods big time, the people are extremely happy," he said, because cops are better due to washed-in silt and the fish are bigger and more plentiful. "Fisherman could take more fish from the river if they know it will be superabundant in one year."

Mosaics also are of interest to scientists studying global change because the areas where there is standing water tend to produce the greenhouse gas methane.

Real-time impacts

Building mosaics of such large regions is one function of the Digital Light Table software, used for examining images too large for traditional viewing programs. JPL's Herb Siegel, chief computational scientist for the Grand Challenge team, explained that the program is a window on the disk, only using workstation memory to pan and zoom through on-screen images.

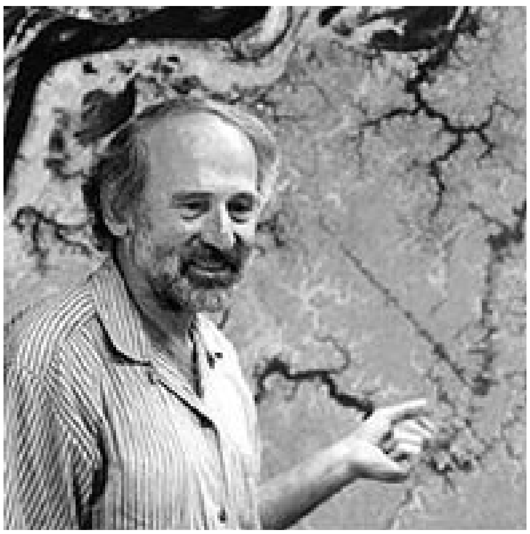

Herb Siegel, JPL

The Digital Light Table will be folded into a Scalable SAR Software Suite (S4) aimed at making SAR more usable. Interferometry processing, drawing on the Scripps algorithms, is one emphasis. Speeding up SAR processing in general will be another prime contribution.

Three years after its two flights, just over 20 percent of SIR-C's data has been processed into image products ordered by researchers. The instrument takes about 15 seconds to record an image, then SIR-C's dedicated supercomputer takes one hour to process that 15 seconds. With NASA HPCC supporting development and access to the latest scalable parallel computers, S4 is nearing 100 seconds per image. A 100 billion flops milestone due in 1999 will reduce that time to 50 seconds, near real-time, according to Craig Miller, JPL software developer.

JPL's Dave Curkendall (left) and Tony Freeman describe the SIR-C antenna superstructure in the background, which is being adapted for the Shuttle Radar Topography Mission, due to fly in 2000.

Freeman stressed that the potential of SAR will never be exploited until it is done in real-time. Speed is especially critical if NASA launches a US SAR satellite. Shi explained that with real-time processing, if you have a big snowfall or disaster, you could analyze the data in days. Similarly, Minster envisions a weekly picture of Southern California tectonic plate evolution. Right now, he bemoaned, they only get data as targets of opportunity. Anticipating these kinds of explorations, JPL is tooling S4 for distributed computing.

"We're moving the computation all over the country via high-performance networks to wherever parallel systems are available," stated Siegel.

"A scientist should be able to say 'I would like to form an interferogram between the ERS-2 passes over Southern California on July 15 and August 31' and initiate a run that finds that data, brings the software to large-scale machines, does the processing, and returns the result in a visual way," Curkendall said. "That whole sequence will transform Earth science data processing from a product to a capability."

"The signals look like white noise because the radar doesn't have a lens. Like the aperture controlling how much light enters a camera lens, SAR gets a wonderful aperture from the motion of the spacecraft. The computer then supplies the lens that focuses the image."- Dave Curkendall, JPL

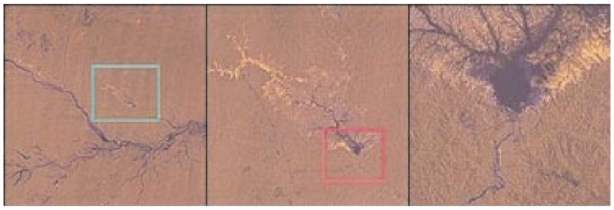

This three-panel view shows successively higher resolution views of the Amazon River basin as mapped by the JERS-1. Flooding is clearly seen below tree canopies (bright yellow) and in open water areas (blue).

Reprinted, with format and style modifications, from Insights magazine, January 98 issue, with permission of the Editor, Judy Conlon. Photographs by unidentified students, approved by Judy Conlon, editor

Learn more about JPL's high-performance computing program at <http://olympic.jpl.nasa.gov/> and <http://www.hq.nasa.gov/hpcc/insights/>. ■

Optimization of a Parallel Ocean General Circulation Model

Ping Wang, Computational Scientist, Daniel S. Katz, Computational Scientist, and Yi Chao, Earth and Space Science Division, Jet Propulsion Laboratory

One of the principal roles of the ocean in the global heat balance is storage of heat, which moderates seasonal extremes and leads to the contrast between ocean and continent temperatures. In the global warming environment, the ocean could delay the onset of the atmospheric warming by absorbing some of the excess heat trapped in the atmosphere.

Ocean modeling plays an important role in both understanding the current climatic conditions and predicting the future climate change. In situ oceanographic instruments provide only sparse measurements over the world ocean. Although remote-sensed data from satellites cover the globe on the time scale of 2~10 days, it only provides information on the ocean surface. Information below the ocean surface has to be obtained from 3-dimensional ocean general circulation models (OGCMs).

OGC models

OGCMs are usually used in the following three types of applications: (1) eddy-resolving integrations at ever increasing resolutions, (2) a large number of model sensitivity experiments, and (3) long-period climate integrations over multiple decades or even centuries. Due to the computational cost of running a 3-dimensional OGCM, these applications are far from having been exhaustively run.

The existing OGCMs do not properly resolve synoptic disturbances (or weather) commonly known as meso-scale eddies in the ocean. Much of the ocean energy is concentrated at a small physical length scale known as the radius of deformation, which varies from about 1 degree (on the order of 100 km) near the equator to 1/10 or even 1/20 degree at high latitudes. This is the scale of intense boundary currents as well as transient eddies, and these phenomena are of considerable importance to the larger scale dynamics. It was not until recent years that eddy-permitting (or eddy-resolving) calculations could be carried out on a basin or global scale. Using the vector supercomputers (e.g., Cray Y-MP) at National Center for Atmospheric Research (NCAR), decade-long ocean model integrations have been carried out at 1/4 degree horizontal resolution [1], which was the first OGCM with performance exceeding 1 billion floating-point-operation-per-second (1 GFLOP/s). With the advance of massively parallel computing technology, decade-long integrations at 1/6 degree resolution have been conducted at Los Alamos National Laboratory (LANL) [2] and Jet Propulsion Laboratory (JPL) [3] on the CM-5 and Cray T3D, respectively. Recently, a short two-year integration at 1/12 degree resolution was made on the T3D at the Pittsburgh Supercomputer Center [4]. Despite the recent progress in eddy-resolving ocean modeling, it is not clear whether 1/6 or even 1/12 degree resolution is sufficient to resolve the ocean eddies and their impact on the large-scale circulation. Some run time estimates of global OGCMS on various advanced computers are given on Table 1.

In addition to running OGCMs at ever increasing resolutions, OGCMs also need to be thoroughly tested against the available observations to establish their validity in describing the real world. This often involves conducting a large number of experiments with different model configurations, and seeking the "best fit" between the model and data. Because of the limited computing resources available to ocean modeling community, it is often a very challenging task to systematically test various combinations of different model parameters.

When studying the Earth's climate, OGCMs have to be integrated over long period of time, on the order of 100s or even 1000s of years. This is mainly because of the long memory of the ocean system, which has a large impact on the other components of the Earth system, e.g., atmosphere, land and ice.

Table 1. Run time estimates of global OGCMS on various advanced computers

Platform/nodes	Sustained Speed (GFLOPS/s)	Grid Size (degree by degree by level)	CPU time (hours/simulated year)
Cray T90/32	15	1/4x1/4x40	10
TMC CM-5/1024	10	1/6x1/6x40	30
Cray T3D/1024	30	1/8x1/8x40	25

Given the limited computing resources available to the ocean modeling community, it is therefore very important to develop an efficient community ocean model, which can be used to conduct the above described applications. It is precisely this objective that motivates the present study. The goal is to select a commonly used OGCM and improve its computational performance, particularly on advanced parallel computers.

Model description

Based upon the above described objective,the most widely used OGCM code was chosen as the base code. The OGCM is based on the Parallel Ocean Program (POP) developed at Los Alamos National Laboratory [2].This ocean model evolved from the Bryan-Cox 3-dimensional primitive equations ocean model [5,6], developed at NOAA Geophysical Fluid Dynamics Laboratory (GFDL), and later known as the Semtner and Chervin model or the Modular Ocean Model (MOM) [7]. Currently, there are hundreds of users within the so-called Bryan-Cox ocean model family, making it the dominant OGCM code in the climate research community.

The OGCM solves the 3-dimensionl primitive equations with the finite difference technique. The equations are separated into barotropic (the vertical mean) and baroclinic (departures from the vertical mean) components. The baroclinic component is 3-dimensional, and uses explicit leapfrog time stepping. It parallelizes very well on massively parallel computers. The barotropic component is 2-dimensional, and solved implicitly. It differs from the original Bryan-Cox formulation in that it removes the rigid-lid approximation and treats the sea surface height as a prognostic variable (i.e., free-surface). The free-surface model is superior to the rigid-lid model because it provides more accurate solution to the governing equations. More importantly, the free-surface model tremendously reduces the global communication otherwise required by the rigid-lid model. Building upon the original ocean model developed at LANL, the new JPL ocean model has significantly optimized the original code, and developed a user-friendly coupling interface with the atmospheric or biogeochemical models.

General optimization strategies

The original POP code was developed in FORTRAN 90 on the Los Alamos CM-2 Connection Machine [2]. During the first half of 1994, the code was ported to the T3D by Cray Research using SHMEM-based message passing. Since the code on the T3D was still time-consuming when large problems were encountered, improving the code performance was required. In order to significantly reduce wallclock time, the code was optimized using single PE optimization techniques [8] and other strategies. The remainder of this section discusses these strategies and the corresponding improvement in performance of the POP code on the T3D.

Memory optimization and arithmetic pipelines

The T3D uses the DEC Alpha EV4 processor with a 151 MHz clock and is capable of 151 MFLOP/s of peak performance. The cache is 8 KB, direct mapped, and has 32 byte cache lines. There are several ways to achieve good performance on the T3D. One is through effective use of cache; another is through effective use of pipelined arithmetic. The POP code uses many two dimensional arrays, which can be inefficient when frequent stride-one addressing is encountered. One improvement is to change to explicit one dimensional addressing. For example, the two dimensional array KMU(IMT,JMT) can be replaced by KMU(IMT*JMT). This type of change can increase performance, both by simplifying index calculations and by making the code easier for the compiler to optimize.

Another useful strategy is to unroll loop statements. The DEC EV4 processor has segmented functional units for floating point multiply and addition. Although a multiply or addition can be issued every clock period, the result is not ready for 6 clock periods. (The divide operation needs 61 clock periods as it is not a pipelined function.) Thus, in order to get top performance from FORTRAN code, you must expose functional unit parallelism to the compiler. Because of these features of the T3D system, the POP code has been optimized by unrolling loop statements if this will expose functional unit parallelism. The number of divide operations also has been minimized by using real variables to store the values resulting from divide operation and moving the divide operation out of loop statements when it is independent of the loop index. After these techniques were applied, the code performed with a significant increase of the total MFLOP/s, as discussed in the results section of this article.

Eliminating IF and WHERE statements by using mask arrays

In the original POP code, many logical IF and WHERE statements were used in distinguishing ocean points from land points. These statements consumed substantial CPU time and reduced pipelining of operations inside loops. The compiler often was not able to efficiently optimize these loops (especially using automatic loop unrolling), since within the IF and WHERE statements some of the computations were quite complex. These IF and WHERE statements were replaced with land/ocean masking arrays, which store the values 1 for ocean points and 0 for land points, and then use multiplies of these values. For example, the statements:

```
 WHERE (KMU.GE.1)
  VUF = SUF
  VVF = SVF
 ELSE
  VUF = 0.0
  SUF = 0.0
 ENDIF
```

were replaced by:

```
CALL SHAD(IMT*JMT,1.0,UMASK,1,SUF,1,0.0,VUF,1)
CALL SHAD(IMT*JMT,1.0,UMASK,1,SVF,1,0.0,VVF,1)
```

where UMASK, VUF, VVF, SUF, SVF, and KMT are all of size (IMT,JMT), and UMSK was defined after the geometry initialization by:

```
DO I=1,IMT*JMT
 IF (KMU(I,1).GE.1) THEN
  UMASK(I,1) = 1.0
 ELSE
  UMASK(I,1) = 0.0
 ENDIF
ENDDO
```

This can be done because KMU is strictly a function of input geometry, and it results in a speed-up of 3.2 times, for IMT = JMT = 100. We have also applied some mathematical formula to eliminate some IF/WHERE statements. For example, in the subroutine STEP the statements

```
if (mix.eq.0) then
   beta  = alpha
   c2dtt = c2*dtt
   c2dtu = c2*dtu
   c2dtp = c2*dtp
 else
   beta  = theta
   c2dtt = dtt
   c2dtu = dtu
   c2dtp = dtp
 endif
```

are replaced by

```
beta=alpha*(1-mix)+theta*mix
c2dtt=c2*dtt*(1-mix)+dtt*mix
c2dtu=c2*dtu*(1-mix)+dtu*mix
c2dtp =c2*dtp*(1-mix)+dtp*mix
```

which will improve the effective use of pipelined arithmetic.

Using optimized libraries

There are several optimized libraries are available on the T3D, such as the SHMEM and BLAS libraries. These libraries have been already optimized to give the best possible T3D performance if the user applies them properly.In the POP code, there are some global sums in the conjugate-gradient routine and in the energy diagnostics routine. Those sums can be performed by using SHMEM library rather using PVM calls. This substitution gives a better performance. There are also many matrices and vectors computations in the POP code which consume a fair portion of the total CPU time. These have been replaced by calling BLAS routines. For example,

```
 DO I=1,IMAX
  X(I)=ALPHA*(Y(I)*Z(I))
ENDDO
```

was replaced with a call to the extended BLAS routine named SHAD:

```
CALL SHAD(IMAX,ALPHA,Y,1,Z,1,0.0,X,1)
```

This can improve performance by a factor of 5 to 10.

<u>Equation of state example</u>

As an example, the routine which computes the equation of state for ocean water was optimized by eliminating the IF/WHERE statements, replacing the FORTRAN 90 array syntax with explicit loop structures, converting nested double-loops and triple-loops to single-loops, and performing explicit loop unrolling. Some mathematical formulae were also applied to eliminate some unneeded work. The overall performance improvement in this routine was 450%. Portability Given the short life cycle of the massively parallel computer, usually on the order of three to five years, we want to emphasize the portability of the ocean model and the associated optimization routines across several computing platforms. Thus far, the JPL ocean modeling effort has mainly been conducted on the T3D. However, the ocean model is now in the process of being converted from T3D to the T3E, and the process to port it to the newly available Hewlett Packard (HP)/Convex SPP-2000, using an MPI version of the POP code, has begun. Preliminary results (Table 2) show that the ocean model running on the SPP-2000 is about three to four times faster than on the T3D. As the memory of the SPP-2000 is four times larger than that of the T3D, this implies one could also run a problem four times larger on the SPP-2000 in the same amount of time.

Table 2. Wallclock time (sec) per time step for T3D, SPP-2000 (MPI codes)

Number of CPUs	SPP-2000 (grid size per CPU)	T3D (grid size per CPU)
16	1.09(64x32x20)	3.20((64x32x20)
64	3.19(128x64x20)	3.34(64x32x20)
256		3.73(64x32x20)

As discussed in the section above, the POP code was ported to the T3D in 1994. The code had previously been rewritten as a message passing code (PVM), and this is the version that was initially ported to the T3D. Those porting this T3D version did not optimize the code beyond trying to choose the fastest code to do identical work as had been done in the previous version. As discussed in the section above, new algorithms and methods to do similar work at a faster rate have now been implemented. The remainder of this section discusses the overall results of this work, and compares them with the initial T3D version of POP.

A test problem was chosen with a local grid size of 37 x 34 x 60 cells. Timings were run for machine sizes from 1 to 256 processors, corresponding to a global grid of up to 592 x 544 x 60. The POP code decompose the grid in blocks in both x and y, and all z data for a given (x,y) is local to one processor. All results shown in this section refer to scaled size problems, where the problem size per processor is fixed and the number of processors is varied.

Scaling performance

Figure 1 shows the run time (in wall clock time) per time step vs. the number of processors, for both the original code and the optimized code. The code running on one processor has been improved by 43%, and as the number of processors involved in the calculation increases, so does the improvement due to the optimization, up to 59% on 256 processors.

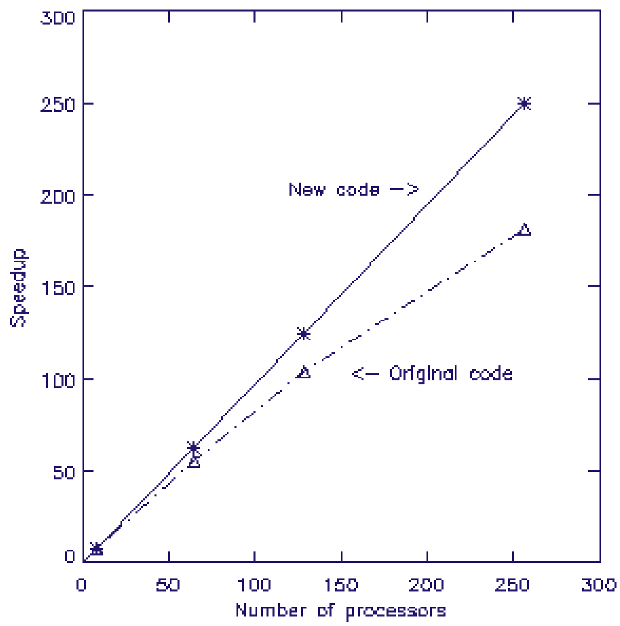

Figure 1. Speed-Up (vs. single processor time) for scaled problems.

Speed-up

It is clear from Figure 1 that the code's scaling has also been improved. This improvement is enumerated in Figure 2, showing the speed-up achieved versus the number of processors used in the calculation. For the ideal parallel code, these two numbers would be equal. The optimized code is performing quite well, in running 250 times faster than the single processor code on 256 processors. The original code, however, clearly had scaling problems, as its speed-up on 256 processors is only 182 times.

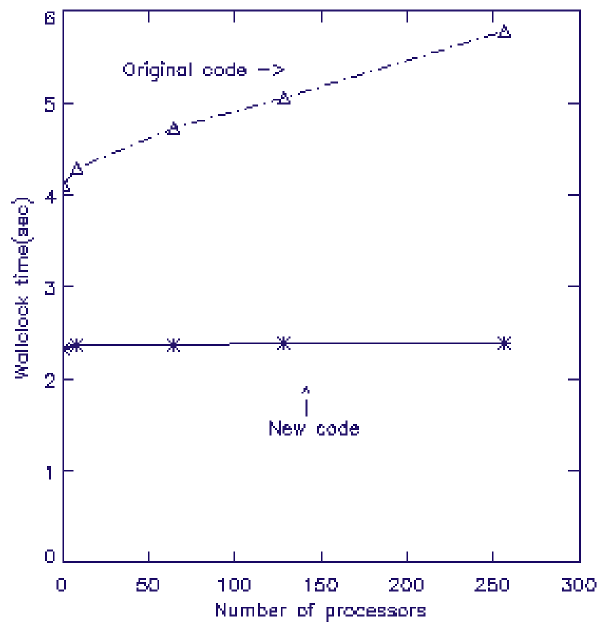

Figure 2. Wallclock time(sec) per time step for scaled problems.

Flop rate

Figure 3 shows the actual performance achieved in terms of computation rate. As mentioned above, the T3D processor has a clock period of 6.6 nanoseconds, corresponding to a clock rate of 151.5 MHz. Since the processor can complete one floating point results per clock period, this is equivalent to a maximum floating point performance of 151.5 MFLOP/s. For 256 processors, this maximum performance is 38.8 GFLOP/s.

It can be observed that the new code is attaining only 10% of the maximum possible performance of the various machine sizes examined. There are three reasons why the code performs at this level. One reason is the ratio of computation to communication in these examples. If a larger local grid size was used, this ratio would increase and the overall performance would also increase. Another reason is poor cache reuse, due to the formulation of the code. It is written in terms of vector-vector routines (replaced by BLAS 1 [9] routines), rather than matrix-vector or matrix-matrix routines (which could be replaced by BLAS 2 [10] or BLAS 3 [11] routines) and uses finite difference routines. Neither of these make enough use of the data (each time it is loaded from memory) to achieve very high performance. (Note that while the data structures being used in the finite-difference routines could be changed to achieve better performance than is now being obtained, this still would be a small fraction of peak performance.) The final reason is that a fairly large number of non-floating point calculations being performed is approximately 13% larger for the optimized code that the original code. This is due to the replacement of IF/WHERE statements with floating point work. The amount of work being done in the code is the same, but some comparison and branch instructions have been replaced by floating point instructions. This is a good example of why floating point performance is a relatively poor method for calculation performance, but as it is considered an important parameter for comparison of unrelated codes, it is included in this paper.

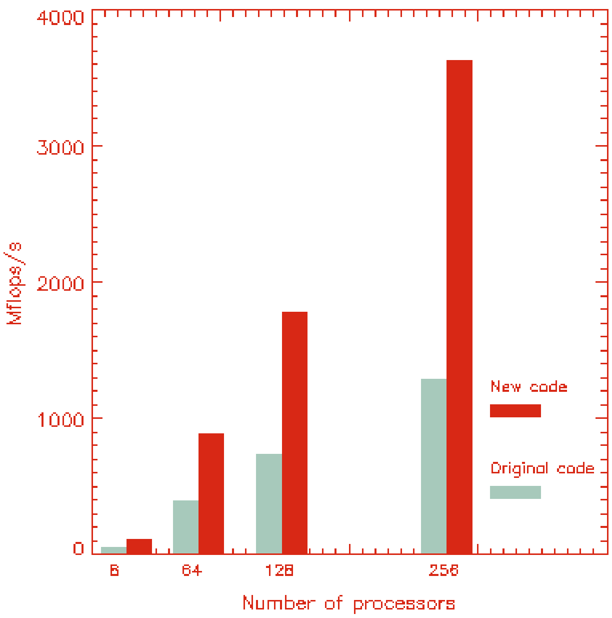

Figure 3. Total floating point performance rates (MFLOP/s) for scaled problems.

Overall simulation results

The test problem size used in the results given above corresponds to either a 1/3 degree global ocean model with 60 depth layers on 512 processors, or a 1/3 degree global model with 30 depth layers on 256 processors. Using the latter model, using the CPU time given as 2.4 sec/ time step, and simulating the circulation of the global ocean for 100 years (or 2.63 million time steps, with each time step being 20 minutes,) would required almost 1800 CPU hours, or about 75 CPU days on the 256 processor T3D. It is possible for this project to use 1.4 CPU days per week of the T3D located at the Jet Propulsion Laboratory, so this large run would take about 53 weeks to complete. This may be contrasted with a similar simulation using the original code, which would take 181 CPU days, or 130 weeks. For large problems such as this, a run which takes one year is possible, but 2.5 years is not. On the HP SPP-2000 which has recently been installed at JPL, the run times is forecast to be about 3 months (once all 256 processors are installed.) The T3E installed at NASA Goddard Space Flight Center should be able to run this problem on 384 processors using about 12 CPU days, though it is currently unclear how much time this would take, since this project doesn't have as much access to that machine.

The test problem may also be used to calculate run times for smaller problems, such as a model of the North Atlantic ocean. A model with 1/6 degree resolution, corresponding to a grid size of 640 x 624 x 45 on 256 processors could be run using approximately the same time. Running this model to simulate one year (at 10 minutes/time step) would require 36 CPU hours. This can be run on the JPL T3D in just over one week. A series of runs can be done to examine how changes in initial conditions will effect the simulation. A simulation of this size may also be compared with measured data, through a repeated process involving variation of physical parameters, to determine correct values for these parameters. The original code would take about 2 1/2 weeks, making these types of analysis much more difficult to perform. Again, on the 256 processor SPP-2000, three to four runs of this size could be done per week, with the optimized code. And on the 384 processor T3E, each run would take about 6 CPU hours.

This 1/6 degree resolution North Atlantic Ocean model (Figure 4) has been run for 30 years of simulated time on the 256-processor T3D, forced with the climatological monthly air-sea fluxes. In comparison with the previous eddy-resolving ocean model simulations [1,2,12,13], this model shows improved Gulf Stream separation off the coast of Cape Hatteras [3]. As the horizontal resolution increases, increasingly fine scale features and the intensification of the currents are found with many of the larger scale features unchanged. It is quite promising that the physical processes responsible for both water mass and eddy formation can be reasonably simulated in models of this class.

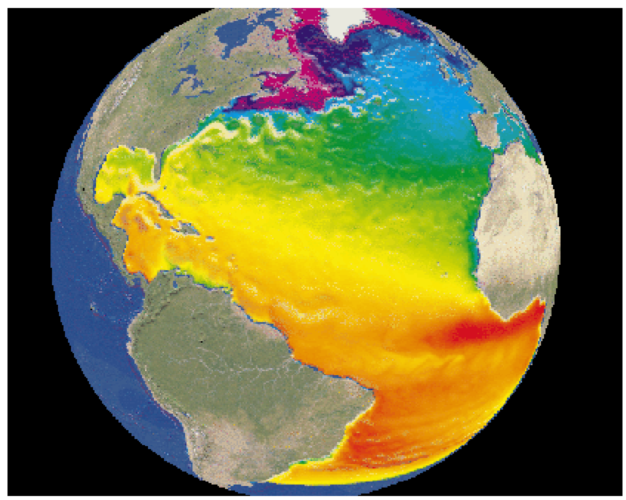

Figure 4. Sea Surface Temperature (in color) and Sea Surface Height Slope (in shade) simulated by a 1/6 degree North Atlantic Ocean model.

Conclusions

Optimization methods which were used on the POP code included unrolling loops, both manually and through the compiler, rewriting code to equivalence multi-dimensional arrays as one-dimensional arrays, simplification of algebra, reductions in the number of divide operations, use of mask arrays to trade fast floating point work for slow comparisons and branches, explicitly rewriting statements in array notation as loops, and using optimized libraries. All of these contributed to the large decrease in time required to solve ocean modeling problems.

The team now anticipates being able to conduct follow-up eddy-resolving integrations at much higher resolutions. The 1/10 and 1/12 degree North Atlantic Ocean calculations have been carried out on the 512-processor CM-5 at LANL and

512-processor T3D at the MSC, respectively. It is quite feasible that JPL can construct a 1/16 degree calculations on the upcoming 384- processor T3E at Goddard Space Flight Center (GSFC) and 256-processor SPP-2000 at JPL. In addition to the OGCM integrations, this OGCM is being coupled with the atmospheric general circulation models developed at UCLA [14,15] on the T3D/E. A biogeochemical component within the OGCM is also being developed, so that the carbon cycle associated with the increasing CO2 and global warming can be addressed.

Acknowledgments and References

The research described in this paper was performed at JPL, California Institute of Technology, under contract to NASA. The Cray Supercomputers used in to produce the results in this paper were provided with funding from the NASA offices of Mission to Planet Earth, Aeronautics, and Space Science. The authors wish to acknowledge Rick Smith, Phil Jones, and Bob Malone for providing the original POP code, David Shirley from Cray Research for doing much of the work in the initial port of this code to the T3D, Alan Stagg from the Army CEWES for doing initial timing studies, Hong Ding from NERSC for initial optimization work, and Martin Lewitt from HP for helping to implement the POP code on the SPP-2000.

1. “Ocean-General Circulation from a Global Eddy-Resolving Model”, J. Geophys. Research Oceans, 97, 5493-5550, 1992.
2. R.D. Smith, J.K. Dukowicz, and R.C. Malone, "Parallel Ocean General Circulation Modeling", Physica D, 60, 38-61, 1992.
3. Y. Chao, A. Gangopadhyay, F.O. Bryan, W.R. Holland, "Modeling the Gulf Stream System: How Far From Reality?", Geophys. Res. Letts., 23, 3155-3158, 1996.
4. R. Bleck, S. Dean, M. O'Keefe, and A. Sawdey, "A Comparison of Data-Parallel and Message-Passing Versions of the Miami Isopycnic Coordinate Ocean Model (MICOM)", Parallel Computing, 21, 1695-1720, 1995.
5. K. Bryan, "Numerical Method for the Study of the World Ocean Circulation", J. Comp. Phy., 4, 1687-1712, 1969.
6. M.D. Cox, "Primitive Equation, 3-Dimensional Model of the Ocean", Group Tech. Report 1, GFDL/NOAA, Princeton, NJ, 1984.
7. R. Pacanowski, R.K. Dixon, and A. Rosati, "Modular Ocean Model User's Guide", GFDL Ocean Group Tech. Report 2, GFDL/NOAA, Princeton, NJ, 1992.
8. J.P. Brooks, "Single PE Optimization Techniques for the CRAY T3D System", Benchmarking Group, Cray Research, Inc., 1995.
9. C..L. Lawson, R.J. Hanson, D. Kincaid, and F.T. Krogh, "Basic Linear Algebra Subprograms for FORTRAN Usage", ACM Trans. Math. Soft., 5, pp. 308-322, 1979.
10. J.J. Dongarra, J. Du Croz, S. Hammarling, and R.J. Hanson, "Extended Set of FORTRAN Basic Linear Algebra Subprograms", ACM Trans. Math. Soft., 14, pp. 1-17, 1988.
11. J.J. Dongarra, J. Du Croz, I.S. Duff, and S. Hammarling, "A Set of Level 3 Basic Linear Algebra Subprograms", ACM Trans. Math. Soft., 16 pp. 1-17, 1990.
12. A.J. Semtner, "Modeling the Ocean General Circulation", Science, 269, 1379-1383, 1995.
13. J.C. McWilliams, "Modeling the Ocean General Circulation", Annual Review of Fluid Mechanics, 28, 215-248, 1996.
14. C.R. Mechoso, C.C. Ma, J.D. Ferrara, J.A. Spahr, and R.W. Moore, "Parallelization and Distribution of a Coupled Atmosphere-Ocean General Circulation Model", Monthly Weather Review, 121, 2062-2076, 1993.
15. L.A. Drummond, J.D. Farrara, C.R. Mechoso, J.A. Spahr, Y. Chao, D.S. Katz, J.Z. Lou, and P. Wang, "Parallel Optimization of An Earth System Model (100 GIGAFLOPS and Beyond?)", 1997 Society for Computer Simulation (SCS) International Conference, Atlanta, Georgia, April, 1997.

This paper was originally published in the SC97: High Performance Networking and Computing Proceedings. SC97 was sponsored by the Association for Computing Machinery (ACM) Special Interest Group on Computer Architecture (SIGARCH) and the IEEE Computer Society Technical Committees on Supercomputing Applications and Computer Architecture.

To learn more about this research, email Ping Wang at <Ping.Wang@jpl.nasa.gov>. Learn more about ACM and the annual supercompting conferences at <http://www-sisn.jpl.nasa.gov/ISSUE45/SC97.html>. Also see NASA Computing technology Featured at SC97" ■

ACM copyright - All individual papers published in the SC97 Proceedings are copyrighted by ACM unless produced as a result of government research.

NSF Approves 29 New Connections to High-Performance Computer Network

The National Science Foundation (NSF) announced February 26 that 29 additional institutions will be connected to the very high performance Backbone Network Service (vBNS), allowing scientists and engineers across the country to collaborate and share powerful computing and information resources. This latest round of connections brings the total number of institutions approved for connections to 92. The vBNS is a crucial player in the president's Next Generation Internet (NGI) and is the initial interconnect for Internet2 member institutions.

"By building an Internet that is faster and more advanced, we can keep the United States at the cutting edge of Internet technology, and explore new applications in distance learning, telemedicine, and scientific research," said President Clinton.

The NSF will make more connections (up to 150 institutions) should the Congress continue to support NSF's role in the NGI appropriation bill directs NSF to use $23 million of the domain name intellectual infrastructure fund toward Next Generation Internet activities. However a preliminary injunction in a pending lawsuit (William Thomas, et al, v. Network Solutions and National Science Foundation) currently prevents NSF from spending this money. For FY 99, NSF has requested another $25 million for NGI activities.

"The vBNS is a facility, like a laboratory or a supercomputer center, that will accelerate science in all disciplines as well as push the limits of networking technology and applications," said George Strawn, director of NSF's Advanced Networking Infrastructure and Research division.

The vBNS, begun in 1995, is an investment of $50 million in a five-year project with MCI Telecommunications Corporation. Connections are evaluated by a peer review process and are approved based on scientific and technical merit. The sophisticated telecommunications network currently runs at 622 million bits per second and is expected to operate at 2.4 gigabits per second (2,400 mbps) by the year 2000. By comparison, the average home modem transmits 28,800 bits per second.The vBNS is expected to always be several steps ahead of commercially available networking. This large capacity allows scientists to collect and share large amounts of data, to collaborate better across large distances, and to run complex equipment from remote sites. The ability to share data and equipment helps scientists studying everything from atoms to galaxies, and to remotely run simulations of science from environment to the beating heart.

Most institutions receive High Performance Connections grants of up to $350,000 from NSF over two years for their connections to offset the cost of linking from their sites to the vBNS backbone. NSF is spending a total of $9,022,859 for this round of connections grants.The institutions approved for high performance connections are:

- Washington University (St. Louis)
- University of Alabama at Birmingham
- University of Alabama in Tuscaloosa
- University of Alabama in Huntsville
- University of Missouri-Columbia
- University of Florida
- Florida State University
- University of Miami
- University of Wyoming
- Washington State University
- Montana State University
- California State Polytechnic University, Pomona
- California State University, San Bernardino
- San Diego State University
- University of California-San Diego
- Wayne State University
- University of Wisconsin-Milwaukee
- Drexel University
- Purdue University
- George Washington University
- Columbia University
- New York University
- University of Massachusetts, Amherst
- Princeton University
- Georgetown University
- University of Idaho
- University of Nebraska-Lincoln
- University of Illinois at Urbana-Champaign
- Cornell University

Excerpted from an NSF press release, #PR 98-13, written by Beth Gaston.

Learn more about NSF at <http://www.nsf.gov/> and NGI at <http://www.ngi.gov/>. ■

Cutting Design Time with VLAB

Pat Kaspar, NASA Research and Education Network, Ames Research Center

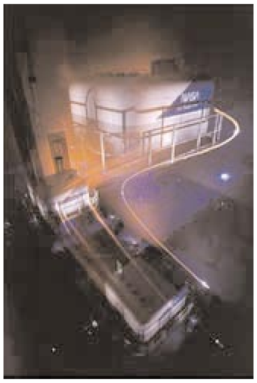

How do you cut costly aerospace design time to keep America on the cutting-edge of a highly competitive global market? NASA addresses that question by aggressively integrating simulation, wind tunnel testing and computational fluid dynamics into the aircraft design cycle, with Ames Research Center (ARC) simulators playing a central role.

In the past, researchers have had to travel to ARC to use the simulation laboratory, SimLab. Now, due to advances in high-performance networking, they can avoid costly, time-consuming travel and participate collaboratively in tests from their own desktops by using "VLAB," a Virtual Laboratory. Remote participation in VLAB simulations is made possible by the high-speed, high-performance NASA Research and Education Network (NREN), part of NASA's High Performance Computing and Communications (HPCC) Program.

Virtual simulations

ARC simulators are high-fidelity, research-oriented facilities that allow industry to repeat design steps with piloted simulations. The ARC Vertical Motion Simulator (VMS) has the world's largest motion basis (40 feet wide by 60 feet tall) and has been used for shuttle, high-speed transport and tiltrotor simulations. The VLAB provides researchers with a collaborative environment in which they can share actual research data.

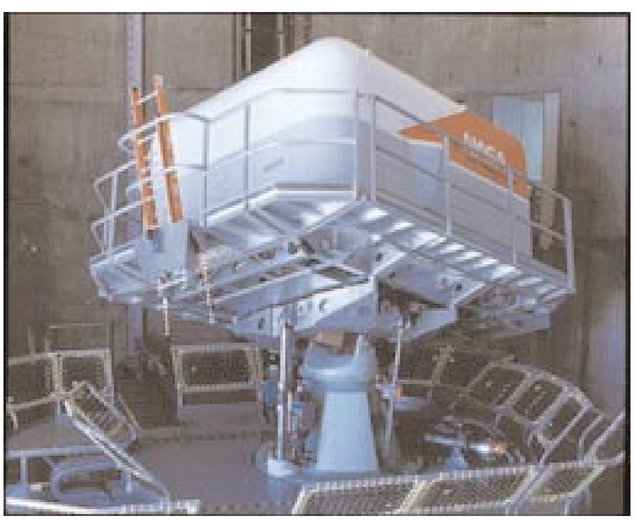

"Our goal," said Chris Sweeney, senior simulation engineer at the VMS, "is to allow you to do the same research you could do if you were in the lab. VLAB was developed in house using a combination of off-the-shelf technology and custom simulation code in use at the VMS. It provides two-way communications, sound, video, data files, real-time animation of the simulator cab in motion and the ability to move around in the lab as if you were there."

The initial successful demonstration of VLAB over an ARC local area network (LAN) proved the concept of a virtual lab.

A proposal was then made to demonstrate VLAB over a wide area network (WAN) connecting ARC with Johnson Space Center (JSC) in Houston, Texas. In order to work in a distributed environment, the VLAB required reliable, high-speed, wide area connectivity to support the client/server connection, along with digital audio, MPEG-2, audio/video conferencing and whiteboarding. That was the challenge presented to NREN in spring 1997. By integrating advanced communications systems and high-performance networking to allow researchers to collaborate in real time, NREN provides a capability which is not possible using today's Internet.

"Engineers have to be able to see what is happening in near real-time so they can react to it," said Ken Freeman, NREN engineering manager. "The technical challenge was to provide a network with latency (delay) low enough to handle the video component, MPEG2, because MPEG picture quality is sensitive to latency and jitter. We needed consistent bandwidth between ARC and JSC."

The first demonstration of the VLAB system via the NREN wide area network communications link was successfully conducted in June 1997 when engineers at JSC participated in VMS Space Shuttle landing simulations.

"With VLAB you can participate as well as just watch," said Steve Cowart, VLAB project lead and manager of the Systems Maintenance and Operational Group at SimLab. "This collaborative tool is especially useful and broadly applicable."

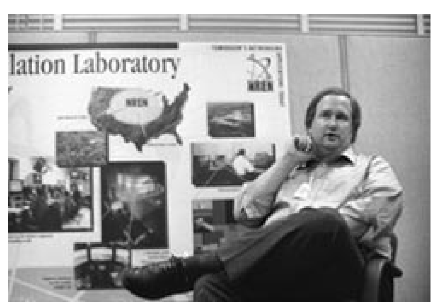

Cowart

Remote research lab

The remote researcher's "lab" is an exact copy of the SimLab Control Room. One monitor shows the same “out-the-window” view that the pilot sees as he lands the shuttle. Another provides two-way videoconferencing, a whiteboard for messages or images, and remote data file access. The central screen is an "active" 3-D model of the Control Room. Researchers can select items and arrange their screens to display, for example, the pilot's Heads Up Display, strip charts, the cab and motion system, a schematic of the shuttle as it lands or all of these. A joy stick allows them to "move" around in the room and even "step back" to see a 3-D version of the entire VMS including the building interior, the control room and the movement of the cab.

Russell Sansom authored the computer interface software that allows full immersion in the simulation. His fascinating 3-D “sticky cursor,” which looks like a green laser beam coming from the bottom of the screen, allows researchers to select various boxes. “It's connected to the sides of the screen by lines that help anchor it”, said Sansom, “and it 'sticks' to objects that can be selected.”

The SimLab team had been working on VLAB for several months prior to the JSC demonstration when Fred Schmitz, Director of Aeronautics, offered support if they would deliver a complete demonstration with JSC in six months. According to Cowart, they used code already developed but added an immersion element so the remote user could share in the experience of being in the SimLab. He explained that flight simulation is virtual reality, and simulations are already remote by nature, even if they are in close proximity physically to the control Room. Because SimLab uses a local network, the team stretched that out over the NREN wide area network to the VLAB client at Johnson Space Center.

“To meet the deadline we worked everything in parallel,” said Cowart, “especially the question of how we were going to get the data back and forth. Our network team looked first within ARC to see who could provide the speed and bandwidth we needed. Our first meeting with the NREN team was very positive. Christine Falsetti, NREN project manager, was already familiar with the project and agreed to partner with us. NREN's readiness to support us was critical in getting the project going.”

“We needed guaranteed bandwidth, as 'guaranteed' as possible, and real-time participation. The researcher at the remote end can't be made to wait to see what's happening. At first we used video technology also for audio, but the software required the researcher to be attached to the computer with a wired headset, a very user-hostile situation. Our solution was to use a pair of products from ASTi and Telex which enabled us to provide sound using remote wireless headsets at both ends.

Cowart explained that simulations play a central role in the aircraft design cycle. The team wanted to use this technology to help reduce the time it takes to develop aircraft.

“It's less expensive to evaluate design steps if the pilot's assessment is part of the process,” he said. “When you combine simulation with computational fluid dynamics and wind tunnel testing, you can make better decisions earlier in the process. VLAB is now a proven research tool that allows remote users to participate and manage the simulation, and to configure their screens the way they desire. Future versions will improve the collaborative feel of VLAB by integrating virtual reality helmets and enhanced input/output elements. You can extrapolate this technology to any kind of product.”

Lessons from the demo

NREN engineers also learned from the demonstration, which clearly showed that MPEG cannot be run over low-performance networks, according to Freeman.

“We learned that this type of video application is very sensitive to competing traffic on the network. We have to be able to guard against latency and jitter in order to provide the quality of service required,” he said.

Remote engineers can't yet give the shuttle a flat tire at the push of a button, but they can request changes in the

parameters of the shuttle landing, using VLAB, which are then carried out by SimLab personnel.

“The JSC demonstration was a huge success and generated a lot of interest,” said Cowart. “We're under tremendous pressure to broaden the application of VLAB today, but we have to balance those pressures with the need to proceed according to plan.”

“Build it and they will come” no longer applies to today's networking world. Scientists and network engineers mustwork together to meet the research challenges of the future. The technologies and methodologies employed in VLAB will be applicable generally to remote access, virtual laboratories offering immersive environments, and interoperable facilities, and will provide broad access to national facilities by U.S. industry. VLAB will provide convenient access to national simulation resources that will enable industry to improve its design process, yielding aeronautical products with a true competitive edge.

NREN has already helped scientists collaborate on echocardiography over great distances, it has made it possible for researchers to operate rovers remotely on Mars and in the Chilean desert, and it proved critical to VLAB's success - collaborative simulation has arrived for NASA.

Reprinted, with format and style modifications, from Insights magazine, September 97 issue, with permission of the Editor, Judy Conlon. Photographs by Judy Conlon and Dominc Hart.

Learn more about VLAB at
<http://www.simlabs.arc.nasa.gov/vlab>, NREN at
<http://www.nren.nasa.gov, and Insights magazine at
<http://www.aero.hq.nasa.gov/hpcc/insights/>. ■

The goal of NASA's many outreach programs is to promote to the general public an understanding of how NASA makes significant contributions to American education systems and to institutions dedicated to improving science literacy. This newsletter provides one vehicle for reporting how applications and hardware used for space science and other NASA research and development can be adapted for use by teachers and their students and by non-NASA organizations.

NASA Conducts Black History Month Internet Chats for Kids

To celebrate Black History Month (February), NASA featured Internet web chats for children with African American scientists and engineers describing their contributions to the missions and goals of the agency. Any youngster with a personal computer with Internet access and web browser software was able to log onto Ames Research Center's (ARC) Quest Learning Technologies Channel World Wide Web site to see a schedule, background information about a mentor, chat instructions, and pre-registration materials. At the scheduled time, the youngster could then go to the chat room and, following the directions, converse with the mentor by typing questions and reading responses. Chat participation was scheduled via a first-come, first-serve pre-registration basis, due to limited capacity by the mentors to answer a large volume of questions. However, chat observation was implemented.

"Each of the scientists and engineers featured during the Black History Month chats was participating in an ongoing program of NASA-related Internet chats intended for children worldwide," said Andrea McCurdy, a project manager at Quest. The chats were conducted with mentors from the Women of NASA, Women of the World, Space Team Online, Mars Team Online, and NeuroLab Online projects.

According to Tish Krieg, also at Quest, the overall mission of these projects is to share NASA with classrooms. "We sponsor online, interactive Internet activities that connect students with NASA people and their work. Research has shown that students learn better by real-life experiences," she said.

Women of NASA chat

The Women of NASA interactive project showcases outstanding women who are enjoying successful careers and how these women balance personal and professional responsibilities.This project was developed to encourage more young women to pursue careers in math, science, and technology.

This chat was scheduled with Kim Hunter, a computer engineer in the Computational Sciences Division of the Information Sciences Directorate at ARC. This division supports major scientific/engineering projects in aeronautics, telerobotics, artificial intelligence, and space systems. Hunter worked on the Wireless Network Experiment, which was performed on both the Space Shuttle Atlantis and the Russian Mir Space Station in March 1996 during the STS-76 mission.

In her background material, Hunter explains that she's always enjoyed math and science and decided to study engineering in college. After high school she attended college, where she earned a Bachelor of Science degree in Electrical Engineering. After college, Hunter served as an officer in the US Air Force working on the development and installation of satellite command and control systems. Upon leaving the service she enrolled in computer science courses and worked part time at ARC. Hunter is mother to twin girls and has many outside interests.

Women of the World chat

An offshoot of the Women of NASA project, the Women of the World project hosts quarterly online chats with the nation's most successful females in a wide range of professions. This project supports students, parents, and schools in an important learning opportunity to dialogue with our nation's most successful female leaders‹typically not accessible to the public‹via the World Wide Web.

This chat was scheduled with Ruth J. Simmons, president of Smith College. Simmons was installed on July 1, 1995, becoming the first African-American woman to head a top-ranked college or university in the United States. Simmons had been part of the university system since 1983, having served at the University of Southern California, Princeton, and Spelman College, then returned to Princeton until her appointment as president at Smith. In 1993, responding to an invitation by the president to review the state of race relations on the Princeton campus, Simmons wrote a report that resulted in a number of initiatives that received widespread attention.

Simmons, the great-great-granddaughter of slaves, was the youngest of twelve children of tenant cotton farmers.

Encouraged by her elementary school teacher, Simmons became gripped by the desire not only to make a difference but also to "achieve something spectacular." Simmons is the recipient of a number of prizes and fellowships, including the German DAAD and a Fulbright Fellowship to France. She was named a CBS Woman of the Year, January 1996; an NBC Nightly News Most Inspiring Woman, August 1996; a Glamour Magazine Woman of the Year, November 1996.

Space Team Online chat

The Space Team Online project provides lots of opportunities to interact with the men and women who are working on the International Space Station and the people who make the space shuttle fly. Youngsters learn about these diverse and exciting careers. This project provides a peek behind the scenes as these experts train astronauts, prepare the shuttle between missions and then launch it, manage the mission, and then bring the shuttle home with a safe landing.

This chat was scheduled with Dr. George Martin, a physician on loan to NASA at Kennedy Space Center (KSC) from the US Air Force. Martin's primary job is to maintain the health of the astronauts and the workers who support the astronauts and to ensure that all employees at KSC, and all visitors who come to tour the Space Center or to watch a launch, enter into a safe healthy environment. Martin also occasionally works in the Launch Control Center as part of the Launch Team.

Martin explains that his career journey was inspired by his wish to be an astronaut and that his first memory is of John Glenn blasting off on the first American orbital mission 35 years ago. Martin was also interested in biology, so he figured he would do research in space biology or space medicine. Because the Air Force helped pay for college, Martin owed them four years of duty, which he spent as a Missile Launch Officer. During that time he earned a master's degree and got accepted to medical school at Ohio State, followed by four years at Johns Hopkins to specialize in emergency medicine.Then he joined the Air Force again and worked in an emergency department for two years. Throughout these years, Martin continued to study space medicine. In 1994 he convinced the Air Force to allow him to work in that field and subsequently came to KSC.

Aero Design Team Online chat

The Aero Design Team Online provides information on why airplanes fly and how they are designed, bringing you into flight simulators and wind tunnels to see NASA employees doing aerodynamic design research. Youngsters hear from the engineers technicians, mechanics and designers working to make tomorrow's planes safer, more efficient, quieter, and faster. Site visitors are able to "look over their shoulders" as these folks operate flight simulations, prepare models for wind tunnel tests, run tests, analyze data, compute fluid dynamic models and more.

This chat was scheduled with Ross Shaw, a wind tunnel test engineer who is essentially responsible for coordinating, planning, scheduling activities, and implementing test objectives defined by the primary investigator. Wind tunnel testing is performed on scaled models to verify the conceptual and preliminary design of an aircraft. Once the optimum configuration is determined, the fabrication phase begins. Flight tests are performed following assembly and systems integration. The flight tests and the wind tunnel tests qualify the aircraft. Wind tunnel testing is continued throughout the life of the aircraft as a means to improve configurations or to help understand problems encountered during the flight operations of the aircraft.

Shaw became an aerospace engineer by accident. In high school one of his assignments as a junior was to decide what he wanted to do, where he wanted to go to college, and in what would he major. Because he had been more interested in basketball than college, he had not really considered this. Because he liked to build model airplanes, boats, and cars, and because he enjoyed math, he thought about aerospace engineering. The class exercise pushed him to learn about the field and what was involved. He enrolled at the University of California at Davis in the aeronautical science and engineering program for five years, followed by graduate school with an emphasis in aircraft structures with the intent to go into aircraft design. After graduation, Shaw interviewed with the Wind Tunnel Facilities Branch to get his foot in the door in the aerospace industry. He was pleasantly surprised to find the wind tunnels at ARC were on the cutting edge of aerospace design system. Thus began his career at a wind tunnel facilities engineer at NASA and later a test engineer.

Mars Team Online chat

The Mars Team Online provides an opportunity to meet with the men and women who work on uncovering the mysteries of Mars. Over the next decade, NASA and other space agencies will be launching many different spacecraft to the red planet. Youngsters can join the enthusiastic people responsible for these missions and learn about their diverse and exciting careers, as well as peek behind the scenes as these folks plan future missions and decide on the appropriate science goals, design and build spacecraft, launch and operate the missions, and study the data which results from the missions.

This chat was scheduled with Dr. Cheick Diarra, manager of the Mars Exploration Education and Public Outreach Office at Jet Propulsion Laboratory. This organization provides curricula resources to teachers and information to the general public regarding Mars missions.

Diarra was born in the republic of Mali, West Africa, where he completed his high school education. Upon graduation he then moved to Paris, France, and attended the University of Pierre and Marie Curie, earning a Bachelor's of Science degree in Mathematics, Physics and Analytical Mechanics. In Washington, D.C., at Howard University, he completed his master's in Aerospace Engineering and a Ph.D. in Mechanical Engineering. Diarra then became a faculty member of the Howard University School of Engineering and taught for five years. In 1988 he joined the Jet Propulsion Laboratory (JPL) as an interplanetary navigator. He currently manages the Mars Exploration Program Education and Public Outreach Office for Pathfinder and Mission Design Team for the future Mars missions.

NeurOn Team Online chat

The NeurOn Team Online provides a unique perspective on NASA's life science research being conducted on the Neurolab Mission. Youngsters are able to follow the behind the scenes action of the men and women involved through regular Field Journals and hear from the space researchers who are anxious for new data about gravitational and space biology. The people preparing the experiments, who live and breathe the excitement of creating space missions will share their thoughts. These reports will offer a glimpse of the day-to-day activities required, including the broad diversity of skills needed as we push on new scientific frontiers.

This chat was scheduled with Janis Davis-Street, a nutritionist in life sciences at Johnson Space Center. Davis-Street is part of a team whose job is determining the nutritional requirements for space flight, that is, finding out what foods and nutrients are important for the astronauts as they venture into space.

Davis-Street grew up in Georgetown, Guyana, which is located on the north eastern coast of South America, and lived there until age 19. She then went to college in Ontario, Canada, and completed graduate school in Alberta, Canada, after which she immigrated to the Houston, Texas area. She feels that her most influential educational experience occurred in her final year of college in an advanced nutrition course. There she investigated and discussed the interactions between nutrition and health, and was interested to see how human behavior could so dramatically affect health.

Black History observance

Black History Month is observed in February in recognition of the birthdays of noted African American pioneers,such as Frederick Douglass, W.E.B. DuBois, Langston Hughes, and Eubie Blake, and institutions, such as the National Association for the Advancement of Colored People and the first Pan African Congress. Additionally, historians recall that the first African American Senator, Hiram Revels, took the oath of office in February, 1870.

Excerpted from NASA press release 98-07 by John Bluck,Information Systems Liaison-Office of External Affairs, at Ames Research Center, information presented in the Ghana Review, written by Yaw Boateng, Professor of Education at Eastern Washinton University, and material on the ARC Quest web site.

Learn more about NASA's observance of Black History Month by visiting the Quest web site's Learning Technologies Channel at <http://quest.arc.nasa.gov/ltc/special/mlk/>. You may access detailed background information on each of the mentors mentioned above, as well as an archived chat “blow-by-blow” session. ■

The goal of NASA's many outreach programs is to promote to the general public an understanding of how NASA makes significant contributions to American education systems and to institutions dedicated to improving science literacy. This newsletter provides one vehicle for reporting how applications and hardware used for space science and other NASA research and development can be adapted for use by teachers and their students and by non-NASA organizations.

NASA Hosts Student Robotics Competition

Ames Research Center (ARC) hosted a student robotics competition featuring students from six San Francisco Bay Area high schools on Tuesday on February 24. Nearly 170 high school students competed against each other from 6-8:30PM in ARC's historic Hangar One as part of the seventh annual For Inspiration and Recognition of Science and Technology (FIRST) Robotics Competition. The contest concluded with FIRST's National Robotics Competition on April 2-4 at Disney's Epcot Center in Orlando, Florida, featuring some 200 student-built robots.

ARC sponsored the Bay Area students in order to stimulate their interest in the fields of engineering, science and technology. The six schools participating in the contest were Aptos High School, Gunn High School, Los Altos High School, Monta Vista High School, Palo Alto Senior High School and Woodside High School.

“NASA is proud to sponsor 170 students in this exciting and educational robotics competition,” said Michael Goldman, ARC's project advisor for the six schools. “We hope the students who participated in this intense, hands-on robotics competition will eventually become NASA engineers.”

What is FIRST?

The FIRST Competition is a national engineering contest which teams up engineers from businesses and industry with high school students. The goal is to motivate youth to pursue careers in science and engineering. However, participating students need not be primarily interested in these fields. For example, a student who likes writing may act as the team's publicist and handle press relations, a student with vocational training can help build the robot, a student strong in math may calculate the required geometry for the robot, a student who enjoys computers can develop a web page for the team, and a student interested in art can design the teamís logo and robot aesthetics. In short, every ability can be applied to the project. Faculty members serve as supervisors and coaches.

Each team starts off with the same standard kit of parts. In six intense weeks, each team brainstorms, designs, constructs, and tests their robot. While the design doesn't need to be elegant, the construction doesn't need to be production quality, and the operation doen't need to be the smoothest, the robot must move, pick up or move another object, and probably lift that object.

Students learn about the design process, safety procedures for power tools, and teamwork during these activities. The effort also involves such diverse aspects as financing the project, coordinating logistics, arranging press coverage for the team, and documenting the impact the team has made in the community. When completed, the robots are shipped to the appropriate regional competition. The teams then compete in a spirited, no-holds-barred tournament complete with referees, cheerleaders, and time clocks.

The competition

Each year, the Competition is different, so returning teams always have a new challenge. Competition details are kept secret until the unveiling at the Kick-Off Workshop, providing a high level of excitement. In addition to the initial workshop, held in New Hampshire, there were five regional events:

- Houston Regional, Space Center Houston, Houston, Texas.
- Motorola Midwest Regional, William Rainey Harper College, Chicago, Ilinois.
- New England Regional, New Hampshire College, Manchester, New Hampshire.
- Johnson & Johnson Mid-Atlantic Regional, Rutgers University, New Brunswick, New Jersey.
- Great Lakes Regional, Eastern Michigan University, Detroit, Michigan.

The National Championship was held April 2-4 at Walt Disney World's EPCOT Center in Orlando, Florida.

Each year the winning team is awarded a Design Innovation Scholarship to Worchester Polytechnic Institute (WPI), a four-year, full-tuition scholarship. The winning team then chooses a student from among its members to receive the award.

“WPI has participated in the FIRST Competition as a sponsor since the competition was founded. We believe strongly in its goals and methods, and it has been a privilege to support the competition with the WPI Design Innovation Scholarship,” said WPI President, Edward A. Parrish

Excerpted from NASA press release 98-09, written by Michael Mewhinney

Learn more about the robot competition, the FIRST organization at <http://www.usfirst.org>, and WPIat <http://www.wpi.edu/>. ■